

**WHAT IS CLAIMED IS:**

1           1. A method comprising:  
2 identifying a plurality of secondary nodes to which an update to data is sent, wherein  
3           at least one secondary node of the plurality of secondary nodes inserts the update in a  
4           respective log of updates to a respective copy of the data; and  
5 sending a notification to each of the plurality of secondary nodes when all of the plurality of  
6           secondary nodes have acknowledged the update.

1           2. The method of claim 1 wherein  
2 each secondary node of the at least one secondary node clears the update from the respective  
3           log of updates in response to receiving the notification.

1           3. The method of claim 2 wherein  
2 clearing the update from the respective log comprises updating a start-of-log pointer in the  
3           respective log.

1           4. The method of claim 2 wherein  
2 the clearing the update from the respective log comprises updating a pointer to a location in  
3           the respective log, wherein  
4           the pointer points to the location if the location contains a next update to clear.

1           5. The method of claim 1 further comprising:  
2 determining that a location of a next update in a first respective log of updates to a first  
3           respective copy of the data at a first secondary node of the secondary nodes differs  
4           from a corresponding location of the next update in a second respective log of updates  
5           to a second respective copy of the data at a second secondary node of the secondary  
6           nodes; and  
7 identifying a set of updates in the first respective log, wherein  
8           each update of the set of updates is not in the second respective log; and  
9 synchronizing the first respective copy and the second respective copy by applying the set of  
10          updates to the second respective copy.

1           6. The method of claim 1 wherein  
2 the determining occurs when a primary node maintaining the data fails.

1           7. The method of claim 1 further comprising:  
2     setting a sent indicator for the update for one of the plurality of secondary nodes when the  
3     update is sent to the one secondary node.

1           8. The method of claim 7 further comprising:  
2     setting a received indicator for the update for the one secondary node when an  
3     acknowledgement of the update is received from the one secondary node.

1           9. The method of claim 8 wherein  
2     the sending the notification to each of the plurality of secondary nodes comprises determining  
3     that a respective sent indicator and a respective received indicator for the update are  
4     set for each of the plurality of secondary nodes.

1           10. A system comprising:  
2     identifying means for identifying a plurality of secondary nodes to which an update to data is  
3     sent, wherein  
4     at least one secondary node of the plurality of secondary nodes inserts the update in a  
5     respective log of updates to a respective copy of the data; and  
6     sending means for sending a notification to each of the plurality of secondary nodes when all  
7     of the plurality of secondary nodes have acknowledged the update.

1           11. The system of claim 10 further comprising:  
2     clearing means for clearing the update from the respective log of updates in response to  
3     receiving the notification.

1           12. The system of claim 10 further comprising:  
2     determining means for determining that a location of a next update in a first respective log of  
3     updates to a first respective copy of the data at a first secondary node of the secondary  
4     nodes differs from a corresponding location of the next update in a second respective  
5     log of updates to a second respective copy of the data at a second secondary node of  
6     the secondary nodes; and  
7     second identifying means for identifying a set of updates in the first respective log, wherein  
8     each update of the set of updates is not in the second respective log; and  
9     synchronizing means for synchronizing the first respective copy and the second respective  
10    copy by applying the set of updates to the second respective copy.

1           13. A system comprising:  
 2    an identifying module to identify a plurality of secondary nodes to which an update to data is  
 3           sent, wherein  
 4           at least one secondary node of the plurality of secondary nodes inserts the update in a  
 5           respective log of updates to a respective copy of the data; and  
 6    a sending module to send a notification to each of the plurality of secondary nodes when all  
 7           of the plurality of secondary nodes have acknowledged the update.

1           14. The system of claim 13 further comprising:  
 2    a clearing module to clear the update from the respective log of updates in response to  
 3           receiving the notification.

1           15. The system of claim 14 wherein  
 2    the clearing module further comprises  
 3           an updating module to update a start-of-log pointer in the respective log.

1           16. The system of claim 14 wherein  
 2    the clearing module further comprises  
 3           an updating module to update a pointer to a location in the respective log, wherein  
 4           the pointer points to the location if the location contains a next update to clear.

1           17. The system of claim 13 further comprising:  
 2    a determining module to determine that a location of a next update in a first respective log of  
 3           updates to a first respective copy of the data at a first secondary node of the secondary  
 4           nodes differs from a corresponding location of the next update in a second respective  
 5           log of updates to a second respective copy of the data at a second secondary node of  
 6           the secondary nodes; and  
 7    a second identifying module to identify a set of updates in the first respective log, wherein  
 8           each update of the set of updates is not in the second respective log; and  
 9    a synchronizing module to synchronize the first respective copy and the second respective  
 10           copy by applying the set of updates to the second respective copy.

1           18. A computer-readable medium comprising:  
 2    identifying instructions to identify a plurality of secondary nodes to which an update to data  
 3           is sent, wherein

at least one secondary node of the plurality of secondary nodes inserts the update in a respective log of updates to a respective copy of the data; and sending instructions to send a notification to each of the plurality of secondary nodes when all of the plurality of secondary nodes have acknowledged the update.

19. The computer-readable medium of claim 18 further comprising: clearing instructions to clear the update from the respective log of updates in response to receiving the notification.

20. The computer-readable medium of claim 19 wherein the clearing instructions further comprise updating instructions to update a start-of-log pointer in the respective log.

21. The computer-readable medium of claim 19 wherein the clearing instructions further comprise updating instructions to update a pointer to a location in the respective log, wherein the pointer points to the location if the location contains a next update to clear.

22. The computer-readable medium of claim 18 further comprising: determining instructions to determine that a location of a next update in a first respective log of updates to a first respective copy of the data at a first secondary node of the secondary nodes differs from a corresponding location of the next update in a second respective log of updates to a second respective copy of the data at a second secondary node of the secondary nodes; and second identifying instructions to identify a set of updates in the first respective log, wherein each update of the set of updates is not in the second respective log; and synchronizing instructions to synchronize the first respective copy and the second respective copy by applying the set of updates to the second respective copy.

23. A computer system comprising: a processor for executing instructions, and a memory to store the instructions, wherein the instructions comprise identifying instructions to identify a plurality of secondary nodes to which an update to data is sent, wherein

at least one secondary node of the plurality of secondary nodes inserts the  
update in a respective log of updates to a respective copy of the data;  
and  
sending instructions to send a notification to each of the plurality of secondary nodes  
when all of the plurality of secondary nodes have acknowledged the update.

24. The computer system of claim 23 wherein  
the instructions further comprise:  
clearing instructions to clear the update from the respective log of updates in response  
to receiving the notification.

25. The computer system of claim 23 wherein  
the instructions further comprise  
determining instructions to determine that a location of a next update in a first  
respective log of updates to a first respective copy of the data at a first  
secondary node of the secondary nodes differs from a corresponding location  
of the next update in a second respective log of updates to a second respective  
copy of the data at a second secondary node of the secondary nodes; and  
second identifying instructions to identify a set of updates in the first respective log,  
wherein  
each update of the set of updates is not in the second respective log; and  
synchronizing instructions to synchronize the first respective copy and the second  
respective copy by applying the set of updates to the second respective copy.